









Interrupts, Continued

- Worth noting that pipelining (very common in current processors) complicates interrupt handling when an interrupt happens, there could be multiple instructions in various stages of execution. What to do?
- "Precise interrupts" are those that happen logically between instructions. Can try to build hardware so that this happens always, or sometimes.
- "Imprecise interrupts" are the other kind. Hardware that generates these may provide some way for software to find out status of instructions that are partially complete. Tanenbaum says this complicates o/s writers' jobs.







I/O Using DMA • Basic idea: Similar to interrupt-driven I/O, but transfer of data to memory done by DMA controller, only one interrupt per block of data. • Complexity versus efficiency tradeoffs similar to interrupt-driven I/O, but may result in fewer interrupts and allow overlap of computation and I/O.























Interrupt-Handler Layer — Processing of I/O Interrupt

- Gets control when requested disk operation finishes and generates interrupt.
- Gets status and data from disk controller, unblocks waiting user process.
 At this point, "call stack" (for user process) contains C library function, system read function, and a device-driver function. We return to the device-driver function and then unwind the stack.



Disks — Hardware Magnetic disks: - Cylinder/head/sector addressing may or may not reflect physical geometry - controller should handle this. - Controller may be able to manage multiple disks, perform overlapping seeks. • RAID (Redundant Array of Inexpensive/Independent Disks): - Basic idea is to replace single disk and disk controller with "array" of disks plus RAID controller. - Two possible payoffs — redundancy and performance (parallelism). - Six "levels" (configurations) defined. Read all about it in textbook if interested. • Optical disks - CD, CD-R, CD-RW, DVD. Okay to skim details!



Disk Arm Scheduling Algorithms

• A little more about hardware: Time to read a block from disk depends on seek time, rotational delay, and data transfer time. First two usually dominate.

• Earlier we said that typical device driver for disk maintains a queue of pending requests (one per disk, if controller is managing more than one). What order to process them in? several "disk arm scheduling algorithms":

- FCFS (first come, first served).
- SSF (shortest seek first).
- Elevator.

How do they compare with regard to ease of implementation, efficiency?























